

# Backend-independent full S3 integration in Apache CloudStack

Daniel Augusto Veronezi Salvador  
<gutoveronezi@apache.org>



# Agenda

- ❖ Introduction
- ❖ Current integration in CloudStack
- ❖ Proposed integration
  - How it works
  - Limitations
  - Future works
  - Demonstration
- ❖ Questions

# Introduction

- ❖ By **backend-independent**, we mean it does not rely on the CloudStack backend to forward/propagate S3 commands to the S3-backend
- ❖ S3 protocol
  - Provides object storage
  - Implements REST

# Current integration in CloudStack

- ❖ Implements limited bucket management
  - Fixed policies
- ❖ Requires changes in the CloudStack backend to support each S3 provider
- ❖ Buckets must be created through CloudStack APIs to be accounted
  - Only the bucket size is accounted
- ❖ Currently implemented for MinIO and Ceph RGW

# Proposed integration

- ❖ Use CloudStack credentials (**apikey** + **secretkey**) to sign the requests to the provider
- ❖ Full bucket management directly through the S3-provider
- ❖ CloudStack credentials are accepted under compliance with the account's role

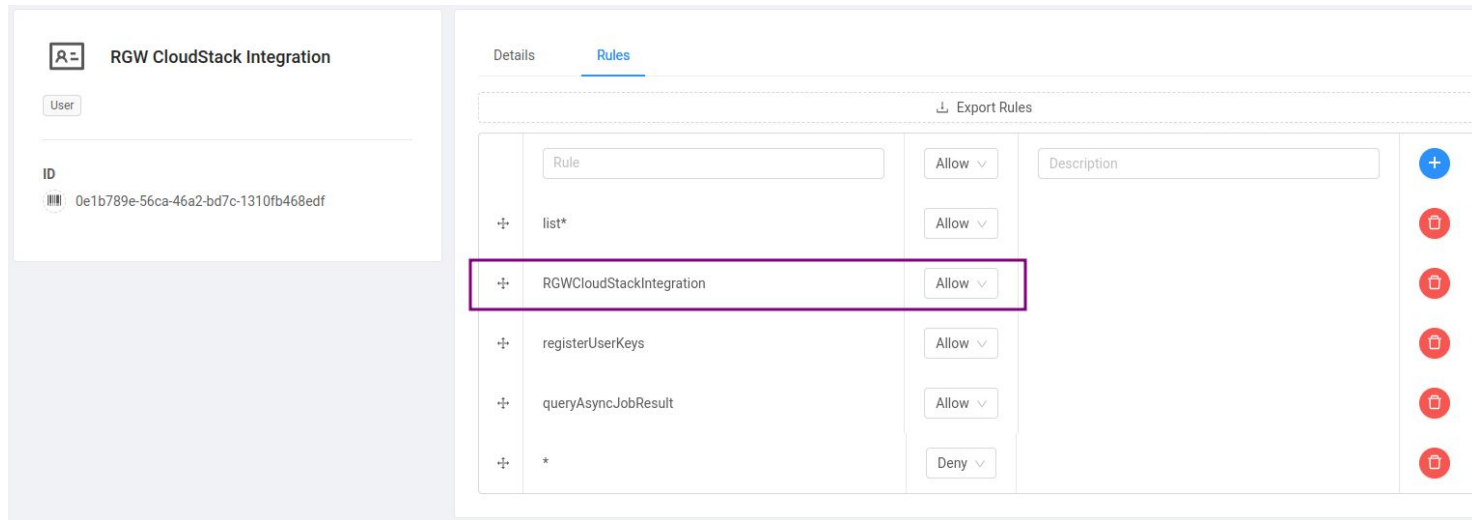
# Proposed integration

- ❖ Introduce CloudStack structures as S3 **principal elements**
  - account
  - domain
  - role
  - roleType
  - user
  
- ❖ GUI for the common features
  - Bucket management
  - Object management
  - Policies management

# Proposed integration

## How it works

- ❖ Define the rule that roles must have to access the S3 provider
  - It can be anything, like **RGWCloudStackIntegration**



The screenshot shows the AWS IAM console interface for a role named "RGW CloudStack Integration". The "Rules" tab is selected, showing a table of permissions. The rule "RGWCloudStackIntegration" is highlighted with a purple box.

Rule	Policy	Description
list*	Allow	
RGWCloudStackIntegration	Allow	
registerUserKeys	Allow	
queryAsyncJobResult	Allow	
*	Deny	

# Proposed integration

## How it works

- ❖ In the Ceph RGW
  - Configure the **CloudStack instance URL**
  - Configure an **apikey** and **secretkey** that has permissions to list all the API key pairs and the rule to access the S3 provider
  - Configure the **rule** that the users must have to access the S3 provider (the same from the role)

```
rgw_cloudstack_rule = RGWCloudStackIntegration
rgw_cloudstack_url = https://cloud.scclouds.com.br/client
rgw_cloudstack_admin_credential_access = Redacted
rgw_cloudstack_admin_credential_secret = Redacted
```



# Proposed integration

## How it works

- ❖ In the CloudStack GUI
  - Configure the URL to the Ceph RGW

# Proposed integration

## Limitations

- ❖ Requires extension in the S3-provider for authenticating with the CloudStack credentials (**apikey** + **secretkey**)
- ❖ Only Ceph RGW was extended so far
  - The changes are not upstream (**yet!**)
- ❖ There is no workflow for accounting the buckets in this proposal (**yet!**)
- ❖ Not all features are addressed in the GUI (**yet!**)
  - But you can do everything the providers support via API
  - The changes are not upstream yet, as well

# Proposed integration

## Future works

- ❖ Flexible rating structures, as for bucket size, operations and so on
  - remember yesterday talk about flexibilizing Usage and Quota?
- ❖ On GUI
  - Multipart upload
  - Bucket lifecycle management
  - Persistent shared bucket listing
  - Object versioning management
  - Wizard for migrating data and structures between S3 providers (e.g.: RGW <-> AWS)
- ❖ Implement the authentication for other Open source S3 providers (like MinIO)?

Let's look at a  
demonstration!

# Questions?

[gutoveronezi@apache.org](mailto:gutoveronezi@apache.org)